Computational Science and Engineering Petascale Initiative at LBL

PI: Alice Koniges

An ASCR Funded ARRA Project

Status: March 16, 2011

- Project Scope: Hire eight or nine post-doctoral researchers for a ~24 month period who will work on key areas supporting DOE's energy mission including energy-related research, ARRA funded projects and Energy Frontier Research Centers (EFRC)
  - Science Goals & Objectives
    - Increase performance, science content and algorithms of energy-related codes running on NERSC systems
    - Provide feedback to NERSC users on application acceleration techniques, new programming models, algorithmic enhancement
    - Encourage and support use of NERSC systems for energy applications
  - Work Authorization Performance Goal:
    - Delivery of computational capability to at least one ERFC by end of the second year.
      - Milestone has been achieved with results from the chemistry EFRC; research continues in this and other areas
  - Budget: $3.125M
    - This budget funds ~9 post-doc positions and their related expenses (workstations and conference travel) for approximately 24 month term assignments.
- Status:
  - 9 positions filled
    - 7 have started working at NERSC, number 8 is due to start April 4, number 9 in May 2011

Petascale Initiative Post-doctoral Hires and Current Projects

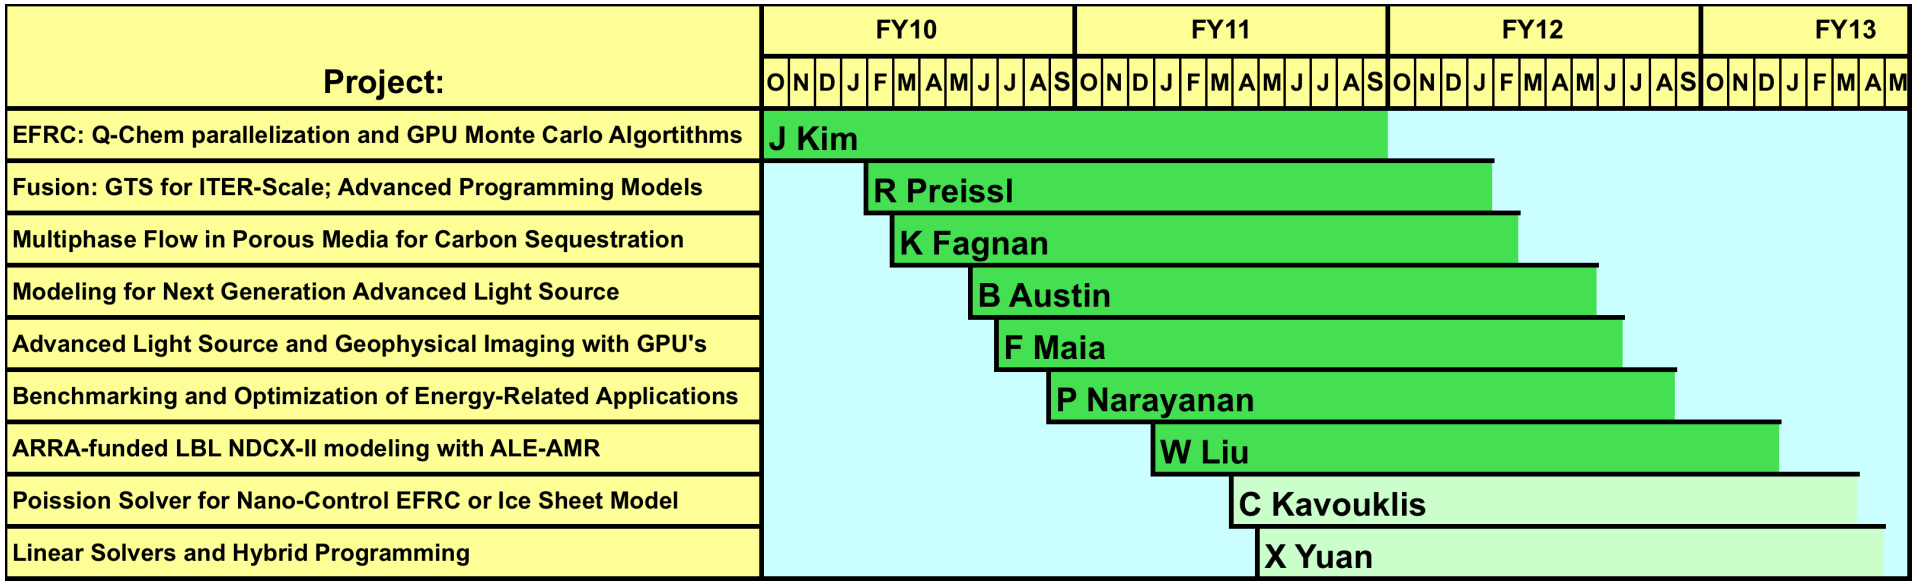

Light green boxes denote projected start date for those who have accepted LBL offers, but not started yet.

Accelerate chemistry codes in order to screen millions of structures

- Codes are used for EFRC Modeling. GPU platforms are introduced.

(1) Parallel Lennard-Jones algorithm
  - CUDA threads within the same block work in parallel to process the same methane – MFI system
  - Total number of independent systems:
    (# CUDA blocks per SM) x (# of SM)
  - Utilize fast GPU memory
  - Can simulate large system size

(2) Embarrassingly parallel algorithm
  - No direct communication between threads (implicit communication through memory access)
  - Total number of independent system:
    (# CUDA threads per block) x
    (# CUDA blocks per SM) x (# of SM)
  - Cannot utilize fast GPU memory
  - Small system size (large DRAM usage)

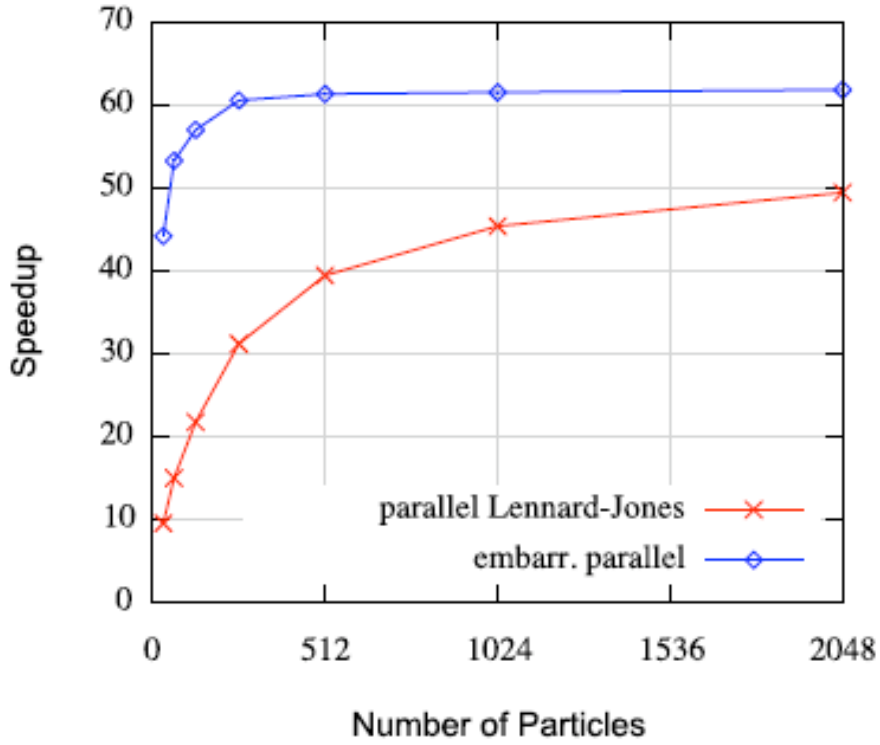

Speedup comparison of the two GPU Monte Carlo algorithms over a single CPU core

Postdoctoral Researcher: J. Kim

Advanced Programming Model Studies using Co-Array Fortran (CAF) in the GTS code

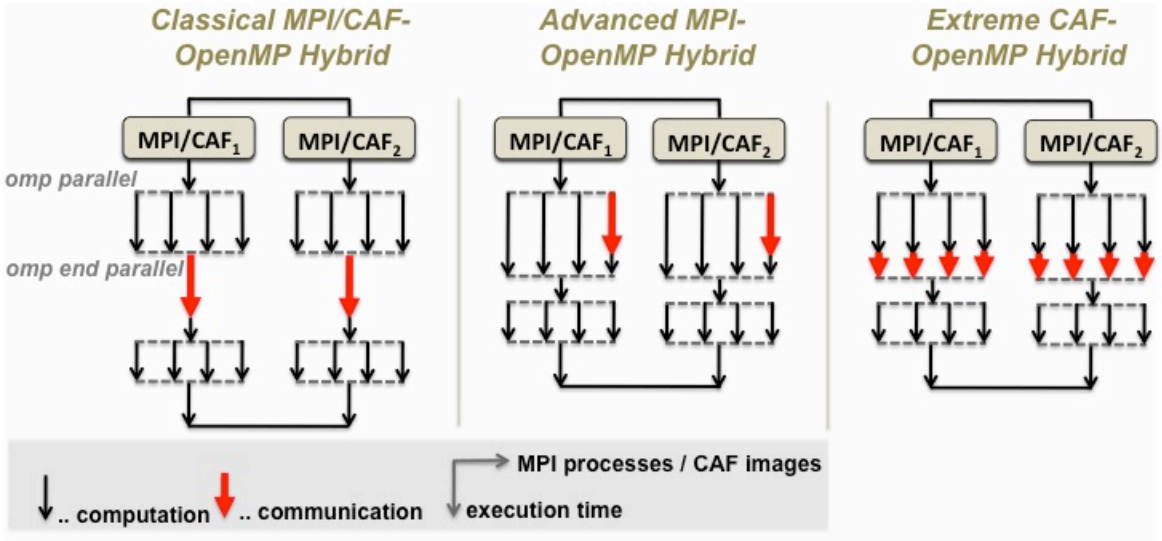

Extend the existing hybrid MPI/OpenMP communication model for better performance and investigate the applicability of new parallel programming models in the communication-intensive part of GTS, a plasma PIC code

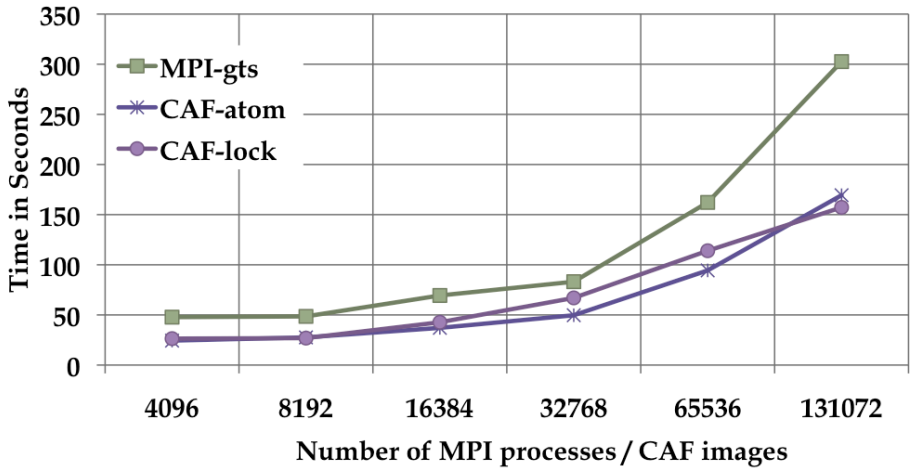

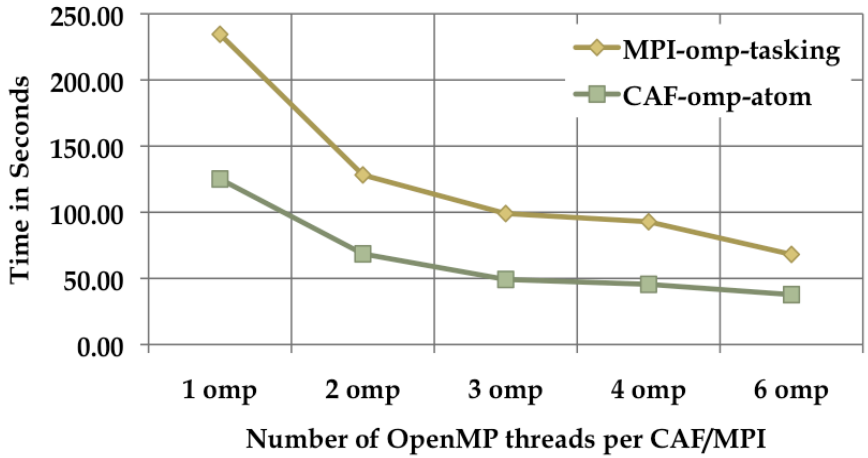

Postdoctoral Researcher: R. Preissl

Performance Optimization of Porous Media Calculations for Carbon Sequestration

- The goal of this work is to enhance the computational and mathematical models used to investigate CO_2 storage in geological formations.

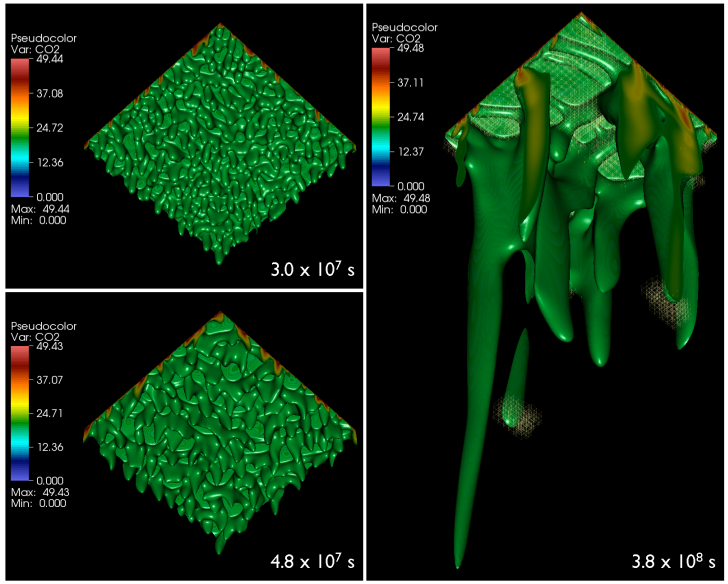

Isosurfaces of CO_2 density in the brine. This image highlights the structure of the 3D fingers.

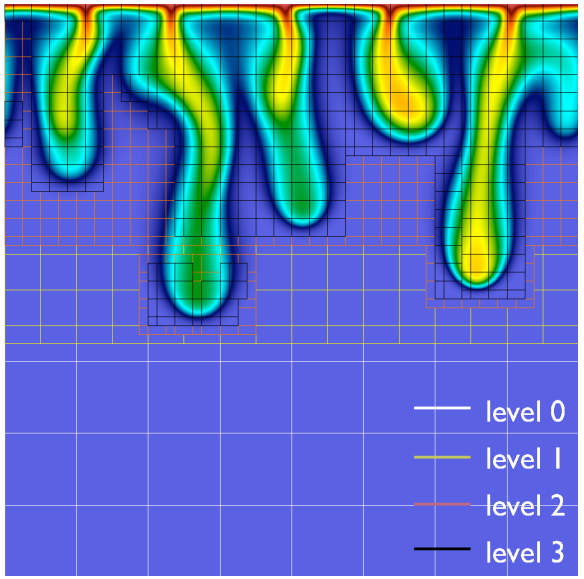

Block-structured adaptive mesh refinement in carbon sequestration calculation. Finest grids are concentrated near edges of the fingers, while coarse grids are used in the fluid, saving computational time.

Postdoctoral Researcher: K. Fagnan

Alternative communication patterns to speed up FFT-based Poisson solver in the Impact-T code

- Impact-T is used for Next Generation Light Source Design

FFT (Computation)

Transpose (Communication)

Original

FFT and transpose performed serially.

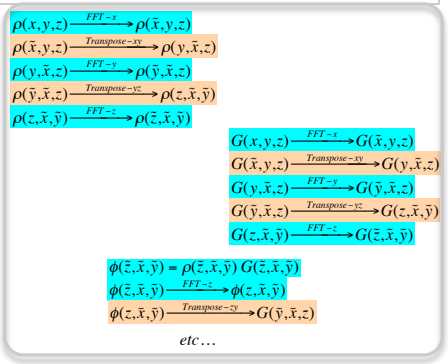

Nonblocking

FFT overlapped with transpose.

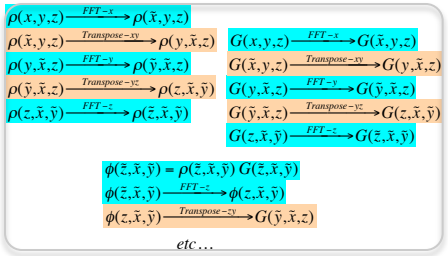

Packed

Transpose messages combined.

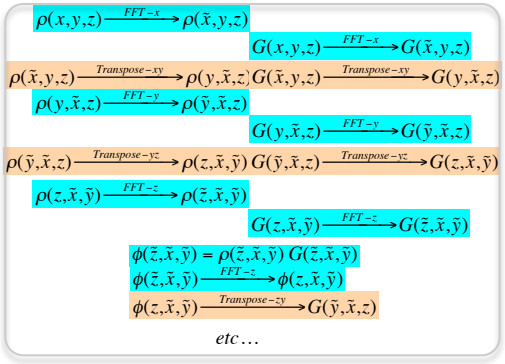

Strong Scaling on Franklin

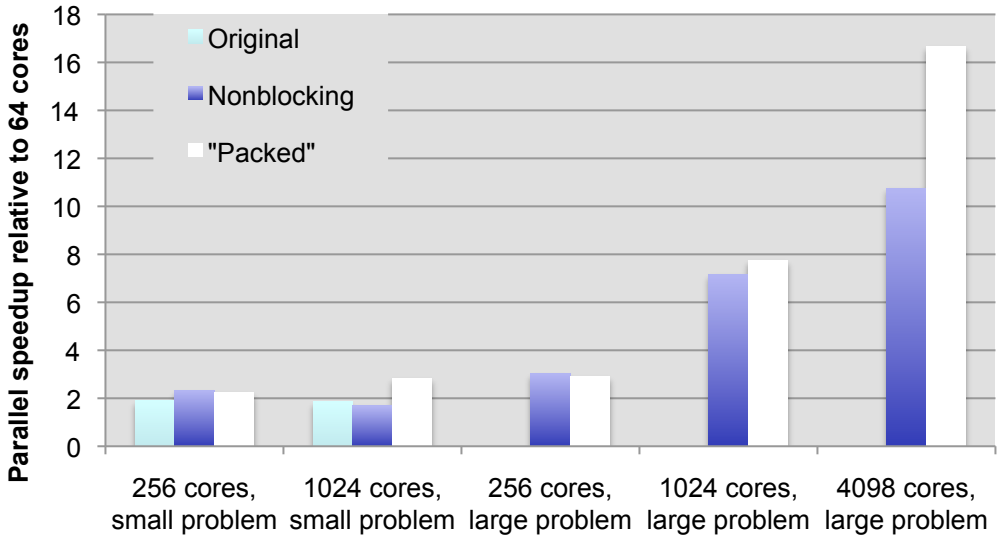

"Packed" algorithm gives 40-60% speedup over original version with 1024 cores on Franklin

Postdoctoral Researcher: B. Austin

Biological Imaging with the Advanced Light Source

- The goal of this work is to advance the analytical tools available to users of the ALS, including the development of real time reconstruction software for the upcoming Nanosurveyor instrument at beamline 9.0.1.

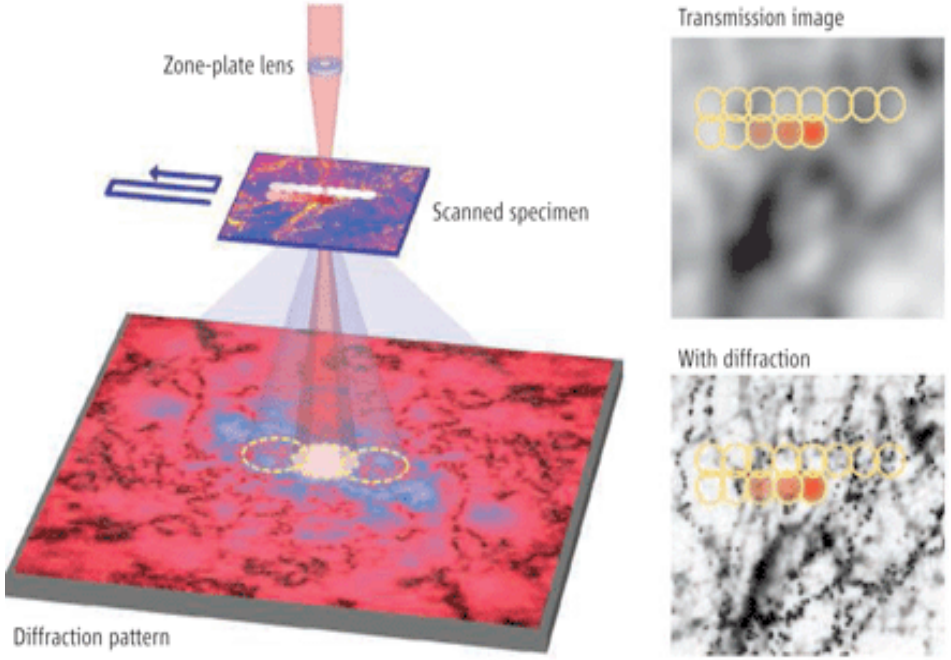

- Matlab implementation completed.
- Serial GPU implementation completed.
- Proof of concept for parallel GPU implementation tested.

Postdoctoral Researcher: F. Maia

Performance Characterization of Magnetic Fusion Applications with Implications for Fusion Co-design

- Characterize application codes for performance on current (petascale) and future (exascale) platforms using NERSC and Cray tools
- Codes: GTS, GYRO/NEO, BOUT++, NIMROD, VORPAL, M3D-C1
- Performance metrics: Scalability, communication, memory and network bandwidth

GTS: Weak scaling

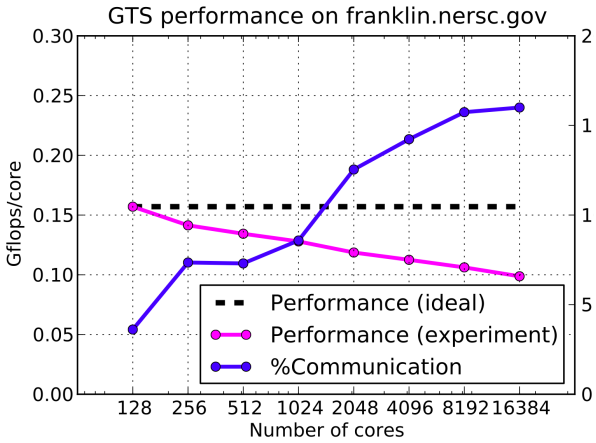

GYRO: Weak scaling

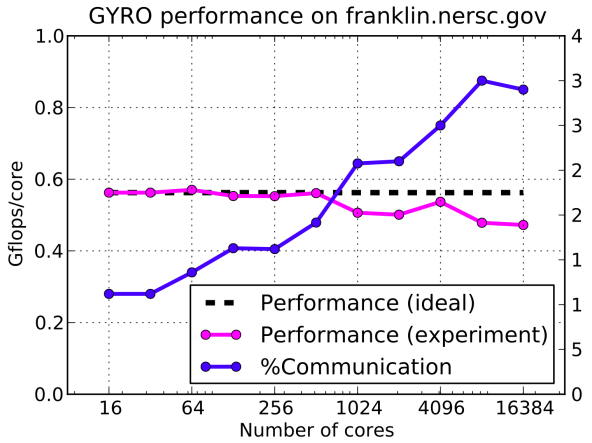

BOUT++: Strong scaling

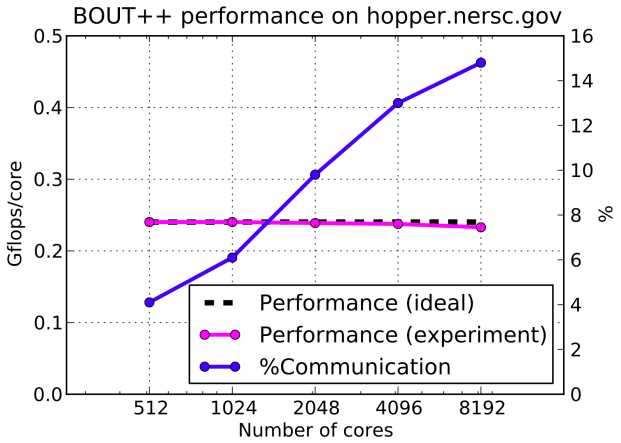

Scaling results from NERSC Franklin and Hopper machines show scalability and role of communication overhead at large concurrency for sample codes.

Postdoctoral Researcher: P. Narayanan

Improving Computational Algorithms and Performance in Warm Dense Matter Modeling

- The goal of this work is to improve the accuracy of NDCX-II modeling codes.
- In particular, we incorporate surface tension effects into the ALE-AMR (Arbitrary Lagrangian Eulerian – Adaptive Mesh Refinement) model.
- NDCX-II is a warm dense matter experiment at LBL with ARRA funding.

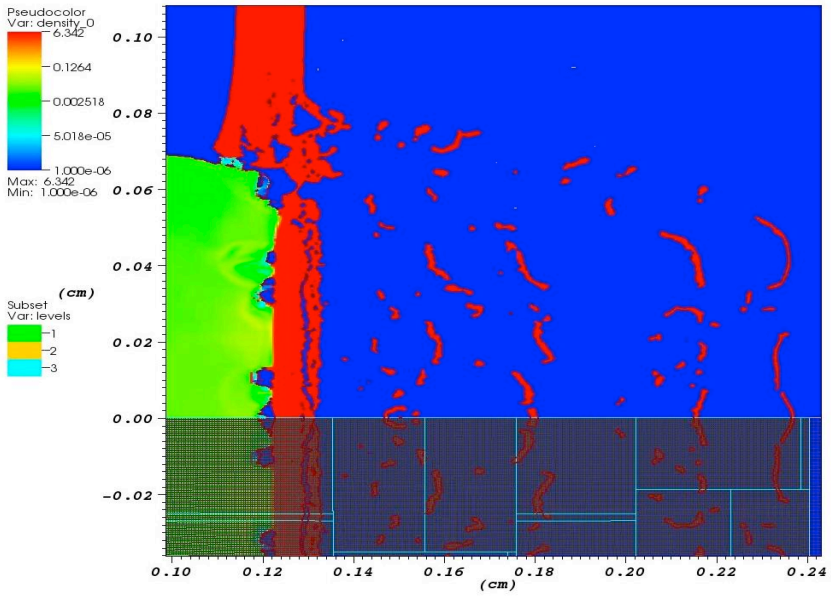

Current modeling of the NDCX-II project without proper surface tension effect. The material breaks up into pieces instead of small droplets.

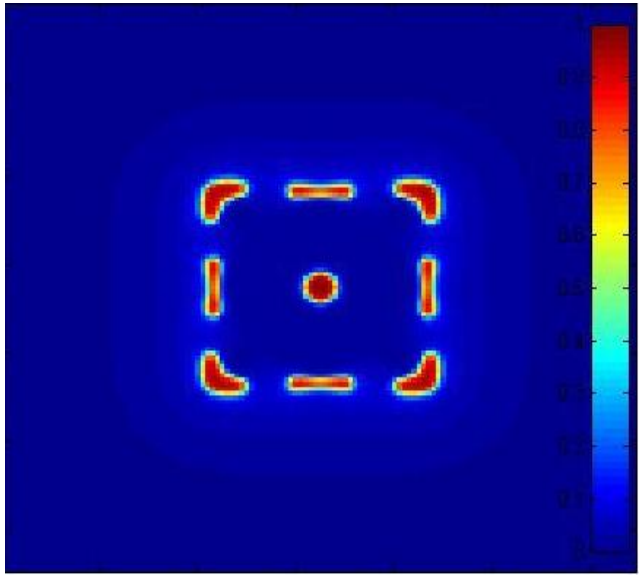

A standalone sample of droplet breakup with a new surface tension model in an expanding flow.

Postdoctoral Researcher: W. Liu

Additional Contributions and Achievements

- Two post-docs (Kim and Preissl) were co-authors and significant contributors to the paper “Application Acceleration on Current and Future Cray Platforms,” winner BEST PAPER, CUG 2010, Edinburgh, May 2010
- Two first author (Fagnan and Narayanan) post-doc paper abstracts were accepted for CUG 2011
- Other papers submitted and/or published include:
  - “GPU Monte Carlo Algorithms for Molecules within a Microporous Framework" - Jihan Kim, Jocelyn Rodgers, Berend Smit, submitted to Caches 2011 workshop.
  - “Acceleration of RI-MP2 Gradient Routine Using Graphics Processing Units" - Jihan Kim, Martin Head-Gordon, submitted to ICCS 2011.
  - “Exploitation of Dynamic Communication Patterns through Static Analysis,” Robert Preissl, B. de Supinski, M. Schulz, D. Quinlan, Dd. Kranzlmüller,and T. Panas, ICPP 2010.
  - “Overlapping Communication with Computation using OpenMP tasks on the GTS magnetic fusion code,” R. Preissl, A. Koniges, S. Ethier, W. Wang, N. Wichmann, Journal of Scientific Programming, 2011.
  - Compressive auto-indexing in femtosecond nanocrystallography, Filipe R.N.C. Maia, Chao Yang and Stefano Marchesini, Ultramicroscopy, November 2010.
  - Compressive phase contrast tomography, F. Maia, A. MacDowell, S. Marchesini, H. A. Padmore, D. Y. Parkinson, A. Schirotzek, and C. Yang, J. Pien, Proc. SPIE 7800, 2010
- More than 10 talks and conference poster presentations were given
- Two post-docs (Kim and Maia) served as teaching assistants for the NERSC-sponsored user course on GPUs, Aug. 2010.
- One post-doc (Preissl) served as a teaching assistant for an SC10 tutorial